| Ref # | Hits | Search Query | DBs | Default Operator | Plurals | Time Stamp |
|---|---|---|---|---|---|---|
| L1 | 17959 | (ddr or (double adj data adj rate)) | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:14 |
| L2 | 803 | (qdr or (quad adj data adj rate)) | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:14 |
| L3 | 18519 | 1 or 2 | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:14 |
| L4 | 88 | 3 with pipelin$ | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:15 |
| L5 | 164608 | (multiplexer or mux or mpx) | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:16 |
| L6 | 39 | 4 and 5 | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:16 |
| L7 | 164685 | (multiplexer or mux or mpx or multi-plexer) | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:17 |
| L8 | 39 | 4 and 7 | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:17 |
| L9 | 1914 | first adj pipeline | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:17 |

| | | | | | | |
|---|---|---|---|---|---|---|
| L10 | 1837 | second adj pipeline | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:17 |
| L11 | 1406 | 9 and 10 | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:17 |
| L12 | 2 | 8 and 11 | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2005/09/15 15:18 |

Terms used **ddr** or **double data rate paragraph pipelines** Found **106,200** of **160,906**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 An overview of the BlueGene/L Supercomputer

NR Adiga, G Almasi, GS Almasi, Y Aridor, R Barik, D Beece, R Bellofatto, G Bhanot, R Bickford, M Blumrich, AA Bright, J Brunheroto, C Caşcaval, J Castaños, W Chan, L Ceze, P Coteus, S Chatterjee, D Chen, G Chiu, TM Cipolla, P Crumley, KM Desai, A Deutsch, T Domany, MB Dombrowa, W Donath, M Eleftheriou, C Erway, J Esch, B Fitch, J Gagliano, A Gara, R Garg, R Germain, ME Giampapa, B Gopalsamy, J Gunnels, M Gupta, F Gustavson, S Hall, RA Haring, D Heidel, P Heidelberger, LM Herger, D Hoenicke, RD Jackson, T Jamal-Eddine, GV Kopcsay, E Krevat, MP Kurhekar, AP Lanzetta, D Lieber, LK Liu, M Lu, M Mendell, A Misra, Y Moatti, L Mok, JE Moreira, BJ Nathanson, M Newton, M Ohmacht, A Oliner, V Pandit, RB Pudota, R Rand, R Regan, B Rubin, A Ruehli, S Rus, RK Sahoo, A Sanomiya, E Schenfeld, M Sharma, E Shmueli, S Singh, P Song, V Srinivasan, BD Steinmacher-Burow, K Strauss, C Surovic, R Swetz, T Takken, RB Tremaine, M Tsao, AR Umamaheshwaran, P Verma, P Vranas, TJC Ward, M Wazlowski, W Barrett, C Engel, B Drehmel, B Hilgart, D Hill, F Kasemkhani, D Krolak, CT Li, T Liebsch, J Marcella, A Muff, A Okomo, M Rouse, A Schram, M Tubbs, G Ulsh, C Wait, J Wittrup, M Bae, K Dockser, L Kissel, MK Seager, JS Vetter, K Yates

November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**

Full text available: pdf(357.61 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper gives an overview of the BlueGene/L Supercomputer. This is a jointly funded research partnership between IBM and the Lawrence Livermore National Laboratory as part of the United States Department of Energy ASCI Advanced Architecture Research Program. Application performance and scaling studies have recently been initiated with partners at a number of academic and government institutions, including the San Diego Supercomputer Center and the California Institute of Technology. This mass ...

2 Tramp: An interpretive associative processor with deductive capabilities

William L. Ash, Edgar H. Sibley

January 1968 **Proceedings of the 1968 23rd ACM national conference**

Full text available: pdf(1.04 MB) Additional Information: full citation, abstract, references, citings, index terms

In recent years; it has become increasingly clear that there is need for a content-addressable computer memory. Larger and larger programs are being written which require a structured data base to operate with any efficiency. Many of these could well benefit by replacing tedious searches with a fast, efficient, "content-addressable" access of the data store. A good example is the "key-word" library search. If one asks for a list of the books written by J. von Neumann ...

3 Frequent value encoding for low power data buses

Jun Yang, Rajiv Gupta, Chuanjun Zhang

July 2004 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 9 Issue 3

Full text available: pdf(1.79 MB) Additional Information: full citation, abstract, references, index terms

Since the I/O pins of a CPU are a significant source of energy consumption, work has been done on developing encoding schemes for reducing switching activity on external buses. Modest reductions in switching can be achieved for data and address buses using a number of general purpose encoding schemes. However, by exploiting the characteristic of memory reference locality, switching activity on the address bus can be reduced by as much as 66%. Till now no characteristic has been identified ...

**Keywords**: I/O pin capacitance, Low power data buses, encoding, internal capacitance, switching

4 A study of performance impact of memory controller features in multi-processor server environment

Chitra Natarajan, Bruce Christenson, Fayé Briggs

June 2004 **Proceedings of the 3rd workshop on Memory performance issues: in conjunction with the 31st international symposium on computer architecture WMPI '04**

Full text available: pdf(316.66 KB) Additional Information: full citation, abstract, references, index terms

With the growing imbalance between processor and memory performance it becomes more and more important to optimize the memory controller features to obtain the maximum possible performance out of the memory subsystem. This paper presents a study of the performance impact of several memory controller features in multi-processor (MP) server environments that use a DDR/DDR2 based memory subsystem. The results from our studies show that significant performance improvements can be obtained by careful ...

**Keywords**: memory controller, memory subsystem, memory transaction scheduling, multi-processors, performance impact, server systems

5 Queue Management in Network Processors

I. Papaefstathiou, T. Orphanoudakis, G. Kornaros, C. Kachris, I. Mavroidis, A. Nikologiannis

March 2005 **Proceedings of the conference on Design, Automation and Test in Europe - Volume 3**

Full text available: pdf(140.78 KB) Additional Information: full citation, abstract

One of the main bottlenecks when designing a network processing system is very often its memory subsystem. This is mainly due to the state-of-the-art network links operating at very high speeds and to the fact that in order to support advanced Quality of Service (QoS), a large number of independent queues is desirable. In this paper we analyze the performance bottlenecks of various data memory managers integrated in typical Network Processing Units (NPUs). We expose the performance limitations o ...

**Keywords**: Network processor, memory management, queue management

6 Hybrid volume and polygon rendering with cube hardware

Kevin Kreeger, Arie Kaufman

July 1999 **Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware**

Full text available: pdf(1.85 MB) Additional Information: full citation, references, citings, index terms

**Keywords**: cube architecture, mixing polygons and volumes, ray casting, run-length-encoding, volume rendering

**7** Unlocking the Performance of the BlueGene/L Supercomputer

George Almasi, Siddhartha Chatterjee, Alan Gara, John Gunnels, Manish Gupta, Amy Henning, Jose E. Moreira, Bob Walkup

November 2004 **Proceedings of the 2004 ACM/IEEE conference on Supercomputing**

Full text available: pdf(170.84 KB) Additional Information: full citation, abstract

The BlueGene/L supercomputer is expected to deliver new levels of application performance by providing a combination of good single-node computational performance and high scalability. To achieve good single-node performance, the BlueGene/L design includes a special dual floating-point unit on each processor and the ability to use two processors per node. BlueGene/L also includes both a torus and a tree network to achieve high scalability. We demonstrate how benchmarks and applications can take ...

**8** Memory Controller Optimizations for Web Servers

Scott Rixner

December 2004 **Proceedings of the 37th annual International Symposium on Microarchitecture**

Full text available: pdf(281.56 KB) Additional Information: full citation, abstract

This paper analyzes memory access scheduling and virtual channels as mechanisms to reduce the latency of main memory accesses by the CPU and peripherals in web servers. Despite the address filtering effects of the CPU's cache hierarchy, there is significant locality and bank parallelism in the DRAM access stream of a web server, which includes traffic from the operating system, application, and peripherals. However, a sequential memory controller leaves much of this locality and parallelism unex ...

**9** Increasing web server throughput with network interface data caching

Hyong-youb Kim, Vijay S. Pai, Scott Rixner

October 2002 **Proceedings of the 10th international conference on Architectural support for programming languages and operating systems**, Volume 30 , 37 , 36 Issue 5 , 10 , 5

Full text available: pdf(1.22 MB) Additional Information: full citation, abstract, references, citings

This paper introduces network interface data caching, a new technique to reduce local interconnect traffic on networking servers by caching frequently-requested content on a programmable network interface. The operating system on the host CPU determines which data to store in the cache and for which packets it should use data from the cache. To facilitate data reuse across multiple packets and connections, the cache only stores application-level response content (such as HTTP data), with applica ...

**10** Session P9: interactive volume rendering: Texture hardware assisted rendering of time-varying volume data

Eric B. Lum, Kwan Liu Ma, John Clyne

October 2001 **Proceedings of the conference on Visualization '01**

Full text available: pdf(11.72 MB) Additional Information: full citation, abstract, references, citings, index terms

Publisher Site

In this paper we present a hardware-assisted rendering technique coupled with a compression scheme for the interactive visual exploration of time-varying scalar volume data. A palette-based decoding technique and an adaptive bit allocation scheme are

developed to fully utilize the texturing capability of a commodity 3-D graphics card. Using a single PC equipped with a modest amount of memory, a texture capable graphics card, and an inexpensive disk array, we are able to render hundreds of time s ...

**Keywords**: PC, compression, high performance computing, out-of-core processing, scientific visualization, texture hardware, time-varying data, transform encoding, volume rendering

**11** Single-Chip MPEG-2 422P@HL CODEC LSI with Multi-Chip Configuration for Large Scale Processing beyond HDTV Level

Hiroe Iwasaki, Jiro Naganuma, Koyo Nitta, Ken Nakamura, Takeshi Yoshitome, Mitsuo Ogura, Yasuyuki Nakajima, Yutaka Tashiro, Takayuki Onishi, Mitsuo Ikeda, Makoto Endo

March 2003 **Proceedings of the conference on Design, Automation and Test in Europe: Designers' Forum - Volume 2 DATE '03**

Full text available: pdf(362.07 KB) Publisher Site

Additional Information: full citation, abstract, index terms

This paper proposes a new architecture for VASA, a single-chip MPEG-2 422P@HL CODEC LSI with multi-chip configuration for large scale processing beyond the HDTV level, and demonstrates its flexibility and usefulness. This architecture consists of triple encoding cores, a decoding core, a multiplexer/de-multiplexer core, and several dedicated application-specific hardware modules with a hierarchical flexible communication scheme for high-performance data transfer. VASA is the worldýs first single ...

**12** Embedded systems: applications, solutions and techniques (EMBS): Assessing the effect of failure severity, coincident failures and usage-profiles on the reliability of embedded control systems

Frederick T. Sheldon, Kshamta Jerath

March 2004 **Proceedings of the 2004 ACM symposium on Applied computing**

Full text available: pdf(327.91 KB) Additional Information: full citation, abstract, references

The increasingly ubiquitous use of embedded systems to manage and control our technologically (ever-increasing) complex lives makes us more vulnerable than ever before. Knowing how reliable such systems are is absolutely necessary especially for safety, mission and infrastructure critical applications. This paper presents a structured compositional modeling method for assessing reliability based on characteristic data and stochastic models. We illustrate this using a classic embedded control sys ...

**Keywords**: design, measurement, performance, reliability

**13** Processor-based system: Unifying memory and processor wrapper architecture in multiprocessor SoC design

Férid Gharsalli, Damien Lyonnard, Samy Meftali, Frédéric Rousseau, Ahmed A. Jerraya

October 2002 **Proceedings of the 15th international symposium on System Synthesis**

Full text available: pdf(682.89 KB) Additional Information: full citation, abstract, references, citings, index terms

In this paper, we present a new methodology for application specific multiprocessor system-on-chip design. This approach facilitates the integration of existing components with the concept of wrapper. Wrappers allow automatic adaptation of physical interfaces to a communication network. We also give a generic architecture to produce these wrappers, either for processors or for other specific components such as memory IP. This approach has successfully been applied on a low-level image processing ...

**Keywords**: embedded memory, memory access, memory wrapper generation, system-on-chip

**14** Polygon rendering on a stream architecture

John D. Owens, William J. Dally, Ujval J. Kapasi, Scott Rixner, Peter Mattson, Ben Mowery
August 2000 **Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware**

Full text available: pdf(161.65 KB) Additional Information: full citation, abstract, references, citings, index terms

The use of a programmable stream architecture in polygon rendering provides a powerful mechanism to address the high performance needs of today's complex scenes as well as the need for flexibility and programmability in the polygon rendering pipeline. We describe how a polygon rendering pipeline maps into data streams and kernels that operate on streams, and how this mapping is used to implement the polgyon rendering pipeline on Imagine, a programmable stream processor. We compare our resul ...

**Keywords**: OpenGL, SIMD, graphics hardware, kernels, media processors, polygon rendering, stream architecture, stream processing, streams

**15** The design and implementation of a new out-of-core sparse cholesky factorization method

Vladimir Rotkin, Sivan Toledo
March 2004 **ACM Transactions on Mathematical Software (TOMS)**, Volume 30 Issue 1

Full text available: pdf(457.74 KB) Additional Information: full citation, abstract, references, index terms, review

We describe a new out-of-core sparse Cholesky factorization method. The new method uses the elimination tree to partition the matrix, an advanced subtree-scheduling algorithm, and both right-looking and left-looking updates. The implementation of the new method is efficient and robust. On a 2 GHz personal computer with 768 MB of main memory, the code can easily factor matrices with factors of up to 48 GB, usually at rates above 1 Gflop/s. For example, the code can factor audikw, currenly the lar ...

**Keywords**: out-of-core

**16** Session P7: unstructured grids and volume rendering: Hardware-software-balanced resampling for the interactive visualization of unstructured grids

Manfred Weiler, Thomas Ertl
October 2001 **Proceedings of the conference on Visualization '01**

Full text available: pdf(752.18 KB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

In this paper we address the problem of interactively resampling unstructured grids. Three algorithms are presented. They all allow adaptive resampling of an unstructured grid on a multiresolution hierarchy of arbitrarily sized cartesian grids according to a varying element size. Two of the algorithms presented take advantage of hardware accelerated polygon rendering and 2D texture mapping. In exploiting new features of modern PC graphics adapters, the first algorithm tries to significantly mini ...

**17** Track 5: supercomputing (part1): ELDORADO

John Feo, David Harper, Simon Kahan, Petr Konecny
May 2005 **Proceedings of the 2nd conference on Computing frontiers**

Full text available: pdf(374.11 KB) Additional Information: full citation, abstract, references, index terms

This paper introduces Eldorado, a third generation multithreaded architecture. Previous Cray multithreaded systems were plagued by unreliable hardware and high costs. Eldorado corrects these problems by using many parts built for other commercial systems. Its compute processor is a 500 MHZ multithreaded processor architecturally similar to the MTA-2 processor; but its interconnection network, I/O subsystem, and service processors are borrowed from other Cray systems. Eldorado retains the program ...

**Keywords**: heterogeneous architectures, multithreaded architectures, multithreaded processing, performance studies

**18** Research papers: stream and sequence mining: Fast and approximate stream mining of quantiles and frequencies using graphics processors

Naga K. Govindaraju, Nikunj Raghuvanshi, Dinesh Manocha

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Full text available: pdf(658.89 KB) Additional Information: full citation, abstract, references

We present algorithms for fast quantile and frequency estimation in large data streams using graphics processors (GPUs). We exploit the high computation power and memory bandwidth of graphics processors and present a new sorting algorithm that performs rasterization operations on the GPUs. We use sorting as the main computational component for histogram approximation and construction of ε-approximate quantile and frequency summaries. Our algorithms for numerical statistics computation on ...

**Keywords**: data streams, frequencies, graphics processors, memory bandwidth, quantiles, sliding windows, sorting

**19** Scalable high-speed prefix matching

Marcel Waldvogel, George Varghese, Jon Turner, Bernhard Plattner

November 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 4

Full text available: pdf(933.02 KB) Additional Information: full citation, abstract, references, citings, index terms

Finding the longest matching prefix from a database of keywords is an old problem with a number of applications, ranging from dictionary searches to advanced memory management to computational geometry. But perhaps today's most frequent best matching prefix lookups occur in the Internet, when forwarding packets from router to router. Internet traffic volume and link speeds are rapidly increasing; at the same time, a growing user population is increasing the size of routing tables against which p ...

**Keywords**: collision resolution, forwarding lookups, high-speed networking

**20** Special issue on knowledge representation

Ronald J. Brachman, Brian C. Smith

February 1980 **ACM SIGART Bulletin**, Issue 70

Full text available: pdf(13.13 MB) Additional Information: full citation, abstract

In the fall of 1978 we decided to produce a special issue of the SIGART Newsletter devoted to a survey of current knowledge representation research. We felt that there were twe useful functions such an issue could serve. First, we hoped to elicit a clear picture of how people working in this subdiscipline understand knowledge representation research, to illuminate the issues on which current research is focused, and to catalogue what

approaches and techniques are currently being developed. Secon ...

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Search Results

Results for "(((ddr <or>(double<near/1>data<near/1>rate))<paragraph>pipelines )<in>..." e-mail
Your search matched **4** of **1235066** documents.
A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

» Search Options

View Session History

New Search

Modify Search

(((ddr <or>(double<near/1>data<near/1>rate))<paragraph>pipelines )<in>metadata)

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Key

| | |
|---|---|
| IEEE JNL | IEEE Journal or Magazine |
| IEE JNL | IEE Journal or Magazine |
| IEEE CNF | IEEE Conference Proceeding |
| IEE CNF | IEE Conference Proceeding |
| IEEE STD | IEEE Standard |

Select Article Information

1. **A 2.5-V, 333-Mb/s/pin, 1-Gbit, double-data-rate synchronous DRAM**
Hongil Yoon; Gi-Won Cha; Changsik Yoo; Nam-Jong Kim; Keum-Yong Kim; Cl Kyu-Nam Lim; Kyuchan Lee; Jun-Young Jeon; Tae Sung Jung; Hongsik Jeong Chung; Kinam Kim; Soo In Cho;
Solid-State Circuits, IEEE Journal of
Volume 34, Issue 11, Nov. 1999 Page(s):1589 - 1599
Digital Object Identifier 10.1109/4.799867
AbstractPlus | References | Full Text: PDF(688 KB) IEEE JNL

2. **A 1.6-Gb/s/pin double data rate SDRAM with wave-pipelined CAS latency**
Sang-Bo Lee; Seong-Jin Jang; Jin-Seok Kwak; Sang-Jun Hwang; Young-Hyur Chil-Gee Lee;
Solid-State Circuits, IEEE Journal of
Volume 40, Issue 1, Jan. 2005 Page(s):223 - 232
Digital Object Identifier 10.1109/JSSC.2004.837983
AbstractPlus | References | Full Text: PDF(1496 KB) IEEE JNL

3. **A High Speed BIST Architecture for DDR-SDRAM Testing**
Sheng-Chih Shen; Hung-Ming Hsu; Yi-Wei Chang; Kuen-Jong Lee;
Memory Technology, Design, and Testing, 2005. MTDT 2005. 2005 IEEE Inter Workshop on
03-05 Aug. 2005 Page(s):52 - 57
Digital Object Identifier 10.1109/MTDT.2005.9
AbstractPlus | Full Text: PDF(184 KB) IEEE CNF

4. **A 1 GHz power efficient single chip multiprocessor system for broadbanc applications**
Santhanam, S.; Allmon, R.; Anne, K.; Blake, R.; Bunger, N.; Campbell, B.; Carl Chen; Cheng, J.; Tuan Do; Dobberpuhl, D.; Ingino, J.; Kidd, D.; Kruckemyer, D Murray, D.; Nishimoto, S.; O'Donnell, L.; Oykher, M.; Panich, M.; Pearce, M.; P Rodriguez, D.; Rogenmoser, R.; Dongwook Suh; Sundaresan, V.; Supnet, E.; Yee, G.; Yiu, G.; Vo, C.; Wen, R.;
VLSI Circuits, 2001. Digest of Technical Papers. 2001 Symposium on
14-16 June 2001 Page(s):107 - 110
Digital Object Identifier 10.1109/VLSIC.2001.934209
AbstractPlus | Full Text: PDF(300 KB) IEEE CNF